Detection of Denial of Service Attacks Against SIP (Session Initiation Protocol) Elements

BACKGROUND AND BRIEF DESCRIPTION OF THE INVENTION

[001] A denial of service attack involves blocking somebody's ability to use some service on a network. Denial-of-Service (DoS) attacks are common across the Internet with many being launched daily at various targets. Many of the attacks involve specially constructed packets or messages designed to either take advantage of flaws in software, or to tie up resources within devices (packet flooding attacks). In some cases, these packet flooding attacks can be directed at devices providing Session Initiation Protocol (SIP) functionality.

[002] SIP is an application level protocol providing multimedia signaling functionality across packet networks. SIP user agents and proxy servers can be used as building blocks to construct IP telephony networks. An example SIP network is depicted in Figure 1. Typical successful SIP message exchanges between a User Agent Client (UAC) and a proxy server for a call initiation are depicted in Figure 2, 3 and 4.

[003] The SIP network architecture can be very flexible, with components distributed throughout IP networks, some trusted (private transport networks), some

not (Internet). Once SIP components are connected to a network that cannot be trusted, the system becomes vulnerable to attacks.

[004] A malicious user can send message floods against SIP elements (as defined in IETF RFC 3261) and render them partially or completely unusable to other users on the network. The present invention provides a solution for detecting malicious INVITE message floods against SIP elements.

[005] In general, for an INVITE flood, the malicious user sends many INVITE messages to a SIP element. The SIP element encounters resource exhaustion when it attempts to keep track of the large number of calls. When legitimate users attempt to make a call, the SIP element is unable to process the messages due to a lack of resources (memory, CPU, etc).

[006] The closest prior art solution to the problem is disclosed in an article by B. Reynolds, and D. Ghosal entitled "*Secure IP Telephony using Multi-layered Protection*", Proceedings of NDSS '03, February 2003 (hereinafter "Reynolds et al). In Reynolds et al, a method is proposed for detection of SIP INVITE message flooding attacks. For each end user, the balance between INVITE and OK messages is used to determine whether the user is under attack. The method uses cumulative sum based change-point detection to analyze when the difference between INVITE and OK is too large.

[007] This prior art solution provides no mechanisms for protecting the infrastructure of the SIP network and the domain of the service provider. An attacker could send a message flood through a proxy server against a non-existent end user. This could result in denial-of-service for all users served by the proxy server. The method described here provides detection mechanisms for the network infrastructure (core and edge proxy servers, etc).

[008] Secondly, this prior art does not take SIP authentication into account. For many systems, some or all of the users will be forced to authenticate themselves to the proxy server when sending INVITE requests. When authentication is introduced the method proposed in this prior art fails.

[009] Finally, the prior art solution uses the balance of SIP INVITE vs. OK messages. The OK message is only sent once the destination user chooses to answer a call. Each user that does not answer their phone results in an imbalance. This could result in additional false positives.

THE PRESENT INVENTION

[010] The present invention provides method and system for detecting DoS (denial of service) attacks against SIP enabled devices. The invention is characterized in that a substantial imbalance between an accounting of SIP INVITE (INV) and SIP 180 Ringing (N_{180}) messages indicates a DoS

attack. This is distinguishable from the prior art, which teaches using an accounting of SIP INVITE and SIP OK messages resulting in more false positives than the present solution.

[011] Preferably, the number (H) of INVITE messages including credentials (INV_c) that are sent from a user client in response to a 407 Authentication Required message from a proxy server are removed from the accounting before the balance is tested. That is, if the equation $INV_0 + INV_c - H = N_{180}$ (where INV_0 is the number of INVITE messages without credentials) is not true within a small margin of error then the presence of a current DoS attack on the proxy server is indicated by the inequality.

DESCRIPTION OF THE DRAWINGS

[012] The above and other objects, advantages and features of the invention will become more apparent when considered with the following specification and accompanying drawings wherein:

[013] Figure 1 depicts an embodiment of SIP network architecture incorporating the invention,

[014] Figure 2 depicts the message handshakes for an unauthenticated SIP call initiation,

[015] Figure 3 depicts the "full" message handshake for a user requiring authentication,

[016] Figure 4 depicts the message handshake for a user with information allowing pre-authentication as described in IETF RFC 2617,

[017] Figure 5 depicts a possible calculated value for detecting SIP INVITE flooding attacks, and

[018] Figure 6 depicts the flow chart for determining whether an INV_c message is part of a full authentication handshake, or a pre-authenticated handshake.

DETAILED DESCRIPTION OF THE INVENTION

[019] The present invention examines the balance between incoming INVITE and outgoing 180 Ringing messages. The 180 message is re-sent by the server once the destination user agent has been identified and has been successfully contacted. The 180 message is sent regardless if the end user answers the call or not and thus identifies a legitimate call.

[020] The present invention differentiates between INVITE messages with authentication credentials given in the "Proxy-Authorization" header field (INV_c) and those without (INV_o). For a system with no authentication enabled (Figure 2), the number of INV_0 messages should be approximately equal to the number of 180 Ringing messages. In a system with authentication enabled for all users (Figure 3) the number of INV_c messages will be approximately equal to the number of 180 messages. For these systems it

is possible to use change-point detection techniques to determine when the two values suddenly are out of balance. The imbalance indicates that a flooding attack is underway.

[021] Detection is more difficult when the system is mixed, with only some users requiring authentication. For the "full" authentication handshake described in Figure 3 there is an INV_o and an INV_c message for a single 180 Ringing message whereas for the handshakes in Figures 2 and 4 there is only one INVITE message per 180.

[022] To identify the "full" authentication handshake the following method may be utilized. A table containing unique call-info values is created. The call-info could consist of call-IDs or digest authentication nonces. When a 407 Authentication Required message is sent from the proxy server to a UAC, the call-info from that message is stored within the table.

[023] For each INV_c message that is received by the proxy, the call-info table is searched. If the call-info from the INV_c message appears in the call-info table, this indicates that the INV_c message is part of a "full" authentication handshake. We label the number of matches as parameter *H*. When a match is found in the table, the entry is then deleted.

[024] The H value can now be used to adjust the balance equation and improve the accuracy of detection. One possible equation for detection is shown in Figure 4.

[025] The same approach can be used for detection of INVITE flooding attacks against User Agents. In this case the invention uses the authentication information found in the 401 Unauthorized messages instead of 407 messages.

[026] The invention allows for analysis of aggregated traffic rather than maintaining statistics per user as is done in Reynolds et al. The problem with per user statistics is the analysis engine may suffer from resource exhaustion due to tracking a large number of users. An aggregated solution does not suffer from this problem.

[027] Secondly, the invention takes systems with authentication enabled into account. It is very likely that at least some of the users will require authentication, so it is not possible to disregard this aspect.

[028] Finally, using the 180 Ringing message rather than the OK message results in less false positives, as answering the call by the destination user is not taken into account in the approach disclosed herein.

[029] While there will be small errors introduced into the system by legitimate calls to incorrect destinations or calls where the user is already on the phone, in these situations the proxy server will receive the INVITE messages, but there will not be any 180 Ringing messages. This should generally occur rarely and should greatly not affect the accuracy of the method.

[030] The ability to detect DoS attacks against SIP-enabled devices is of great value to operators of network services. Efficient DoS detection mechanisms may prove to be value-adding differentiators in the network equipment market. Competitors who add such features to their network equipment may find themselves at an advantage.

[031] While the invention has been described in relation to preferred embodiments of the invention, it will be appreciated that other embodiments, adaptations and modifications of the invention will be apparent to those skilled in the art.